ATTR: An Adaptable and Trustworthy Targeting Recommendation Framework under Real-World Constraints

Sanghyeon Lee†
AX Technology Group
LG Uplus Corp.
Seoul, Republic of Korea
shlee145@lguplus.co.kr

Jinmo Kang†
AX Technology Group
LG Uplus Corp.
Seoul, Republic of Korea
jkang@lguplus.co.kr

Yeonghwan Jeon
AX Technology Group
LG Uplus Corp.
Seoul, Republic of Korea
yeonghwan@lguplus.co.kr

Hyuncheol Jo
AX Technology Group
LG Uplus Corp.
Seoul, Republic of Korea
hcjo@lguplus.co.kr

Byoung-Ki Jeon
AX Technology Group
LG Uplus Corp.
Seoul, Republic of Korea
bkjeon@lguplus.co.kr

ABSTRACT

Evaluating data-driven targeting in enterprise marketing is challenging—analyses are often fragmented, models rarely reusable, and randomized controlled trials (RCTs) typically infeasible. We present ATTR (Adaptable and Trustworthy Targeting Recommendation Framework), a reusable experimentation framework for subscription outbound campaigns. ATTR comprises three layers: modeling for stable target construction, targeting with operator-oriented explainability, and causal evaluation under enterprise constraints. Applied in a subscription campaign, ATTR expanded beyond rule-based pools and surfaced overlooked segments that performed at premium levels. Conversion rates improved substantially, with an average uplift of 23.5 percentage points between the pre-intervention baseline and the post-intervention period. ATTR also surfaced low-purchase users with strong viewing activity who responded well to explanation-based scripts. Causal analysis confirmed these effects: a Difference-in-Differences design estimated a net uplift of 15 percentage points, indicating that improvements were attributable to ATTR rather than external factors. As a modular and reusable framework, ATTR offers a systematic approach to experimentation in subscription marketing and can be extended to other domains requiring trustworthy targeting.

KEYWORDS

Recommender systems, Explainable AI, Causal inference, Trustworthy AI, Marketing applications

†Both authors contributed equally to this research (Co-first authors)

1 Introduction

Outbound subscription marketing depends on targeting quality to drive efficiency and revenue. Rule-based approaches, which rely on explicit purchase histories, cover a narrow group of high-value customers and lose effectiveness over time, leaving most users in low-performing segments.

To overcome these inefficiencies, enterprises have increasingly turned to model-based targeting. Model-based targeting offers broader coverage but faces three barriers in practice. First, analyses are fragmented across teams, preventing end-to-end evaluation. Second, models are rarely reusable and are often tied to single campaigns. Third, RCTs are infeasible in hierarchical workflows where eligibility is predetermined. These barriers make it difficult to identify the true uplift from data-driven targeting.

Prior studies have proposed new predictive or causal models, but few have examined how such methods can be implemented in large-scale operational environments. ATTR addresses this gap through a three-layer design—modeling stable targets, providing operator-facing explanations, and enabling causal evaluation when randomization is impractical.

Applied in subscription marketing, ATTR improved campaign efficiency and uncovered valuable segments overlooked by rule-based approaches. Rather than proposing a new algorithm, ATTR provides an operational framework that integrates modeling, explanation, and causal evaluation for enterprise deployment.

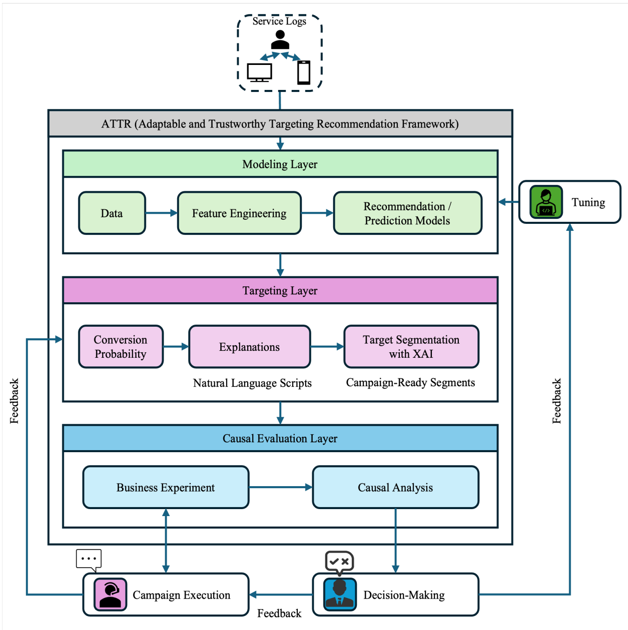

Figure 1. Three-layer modular pipeline of ATTR (Modeling, Targeting, Causal Evaluation), where each component is plug-in and replaceable.

2 Related Work

Research on trustworthy AI has emphasized principles such as transparency, robustness, and fairness, offering conceptual guidance but limited operational use in enterprise workflows [1, 2, 3]. Recent work has called for frameworks that translate these principles into standardized and reusable structures for daily operations [4, 5]. ATTR responds to this need by embedding such concerns directly into a modular experimentation pipeline for subscription marketing.

In recommender systems, explainability research has traditionally centered on end users, aiming to enhance transparency and satisfaction [6, 7]. In enterprise marketing, however, explanations are consumed primarily by operators such as counselors and campaign managers, who require concrete behavioral evidence to support persuasion. Recent studies in operational research have advanced frameworks for operator-oriented XAI [8]. ATTR extends this line by transforming predictive signals into natural-language scripts that counselors can use in practice.

Causal inference has become increasingly important in advertising and recommendation, ensuring that observed uplifts reflect causal effects rather than correlations [9, 10, 11]. While randomized controlled trials remain the gold standard, they are often infeasible in hierarchical workflows. Quasi-experimental methods such as Difference-in-Differences (DID), matching, and Structural Causal Models (SCM) have thus gained prominence [12]. ATTR reflects this evolution by adopting DID as its primary evaluation method while maintaining flexibility to integrate alternative approaches.

In sum, prior studies have advanced principles of trustworthy AI, explainable recommendation, and causal inference largely in isolation. ATTR integrates these strands into a single reusable framework, aligning model reuse, operator-facing explanations, and modular causal evaluation within real-world subscription campaigns. Unlike prior frameworks, ATTR operationalizes trustworthiness by embedding causal validation and explainability into the targeting process, achieving reliability within the system rather than through post-hoc checks.

3 Experimentation Framework (ATTR)

ATTR is organized as a three-layer modular framework—Modeling, Targeting, and Causal Evaluation (Figure 1). Each layer is designed as an independent plug-in module so that algorithms or evaluation strategies can be substituted without compromising the integrity of the pipeline.

3.1 Modeling Layer

The modeling layer ingests service logs—including demographics, subscription records, viewing histories, and purchase transactions—and processes them through feature engineering before passing them to prediction models. The preprocessing pipeline standardizes customer identifiers, normalizes temporal features, and encodes categorical fields such as content genre and device type. This ensures interoperability across campaigns and datasets. In our implementation, we adopted WiDeX, a hybrid wide-and-deep architecture originally developed for IPTV personalization [13]. WiDeX combines wide linear features (e.g., demographic attributes and subscription history) with deep sequential embeddings that capture temporal viewing patterns. This choice was motivated by its robustness under sparse purchase logs and interpretability of feature contributions, which align with the trustworthy design goals of ATTR. Beyond prediction, the modeling layer also outputs explainability signals such as SHAP values. These are not exposed directly to operators but are carried forward for later transformation into interpretable narratives. The modular design allows alternative architectures (e.g., CF or transformer variants) without affecting downstream components.

3.2 Targeting Layer

Predictions generated in the modeling stage are transformed into conversion probabilities and paired with explanations for operator use. These two elements are integrated within the target segmentation module, which constructs campaign-ready customer groups. To operationalize explainability, SHAP attributions were converted into short natural-language templates through a rule-based template engine, producing phrases such as "frequent viewing of action movies in the past week" or "recent interest in premium family content." These templates were further contextualized by customer segment type to generate counselor-ready statements. These cues gave counselors clear openings with customers who had little or no payment history—a setting where traditional persuasion often fails.

During continuous deployment, internal monitoring showed a positive correlation between the frequency of explanation-script usage and conversion outcomes, suggesting that operators gradually adapted to these evidence-based interactions. The targeting layer also supports other explanation or segmentation methods within its plug-in design.

3.3 Causal Evaluation Layer

Campaign-ready segments from the targeting layer flow into the causal evaluation layer, where both execution and validation occur. This layer is composed of two interchangeable modules

3.3.1. Business Experiment Module

This module governs campaign execution under different experimental designs. While randomized A/B testing remains the gold standard, enterprise workflows often predefine eligible customer pools, making randomization infeasible. Accordingly, ATTR accommodates quasi-experimental designs such as Difference-in-Differences (DID) — which was the primary evaluation method used in this study — as well as optional agent-based simulations for offline validation. In the agent-based mode, simulated agents built from historical behaviors enable offline testing of different targeting policies, supporting safer deployment when live trials are infeasible. Although not applied in the current study, this feature illustrates the framework's extensibility beyond quasi-experiments.

3.3.2. Causal Analysis Module

The causal analysis module validates campaign outcomes using inference techniques suited to non-randomized contexts. In our study, a Difference-in-Differences (DID) design estimated a 15 percentage-point net uplift attributable to model-driven targeting — distinguished from a raw improvement of 23.5%p observed over the pre-intervention baseline (see Section 5). Depending on data availability and identification assumptions, this module can also be replaced with matching, structural causal models (SCM), or double machine learning (DML). Through these three layers, ATTR transforms raw service logs into predictions, explanations, target segments, and credible causal evaluations. Developers can extend or replace the modeling module to ensure reusability; operators can apply explanation-enabled targeting to guide customer interactions; and decision-makers can rely on causal analysis modules for reliable assessments of business impact. This integration of prediction, explanation, and evaluation under constrained operational settings distinguishes ATTR from existing enterprise frameworks that typically address only one of these aspects.

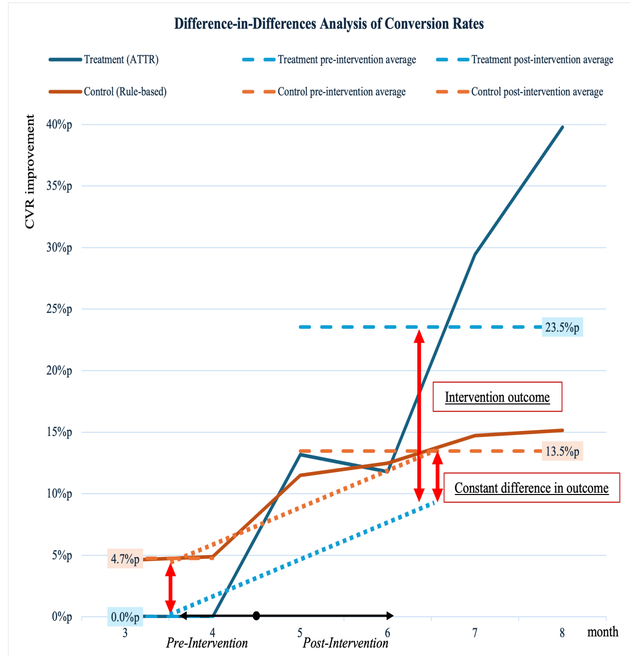

Figure 2. Difference-in-Differences analysis of conversion rates, showing parallel pre-intervention trends and a 15%p intervention effect attributable to ATTR

4 Experimental Setting

ATTR was evaluated over six months using IPTV subscription data. The dataset included demographics, service logs (free and paid viewing), purchase records, and prior marketing responses. Conversion rate was the primary metric, defined as the share of contacted customers who subscribed. We also tracked pool size to assess coverage and revenue per acquisition (RPA) as a financial measure.

The study had two intervention phases. In May, ATTR replaced the existing heuristic targeting applied to the same low-purchase customer segment. Rule-based targeting continued in parallel for frequent purchasers and premium subscribers as part of the standard operation. In June, an explainability component was added: SHAP-based attributions were translated into natural-language scripts and provided to counselors. Monitoring continued through July and August to assess persistence and operator adoption of explanation-based scripts. Baselines were twofold: (1) rule-based targeting of frequent purchasers and premium subscribers, and (2) Heuristic targeting for low-purchase users, based on predefined business rules. The AI group consisted of the latter, ranked by the WiDeX model and selected from the top decile. This setup enabled a clear comparison of ATTR against established practices.

5 Results

5.1 RQ1. Does model-based targeting improve conversion rates compared to rule-based baselines?

The first research question examined whether ATTR improved conversion rates relative to existing marketing practices. In May, when ATTR was first applied to low-purchase customers, conversion rates increased by 13.1 percentage points compared to the baseline segment targeted by heuristic rules. In June, after explanation-based scripts were introduced, conversion rates improved by 11.9 percentage points. Sustained monitoring in July and August revealed even larger uplifts: 29.6 and 39.9 percentage points above baseline, respectively. For confidentiality reasons, absolute conversion rates cannot be disclosed. All reported values therefore denote percentage-point (pp) improvements relative to each segment’s baseline performance rather than absolute figures. These values represent raw uplifts over the pre-intervention period. A Difference-in-Differences (DID) analysis later confirmed that approximately 15 percentage points of this total uplift were causally attributable to ATTR (see Section 5.3). In relative terms, these gains correspond to two to four times the conversion levels of traditional rule-based approaches, effectively closing the gap with the historically high-performing premium group (Table 1). ATTR’s ability to generate sustained improvements across four months suggests operational stability rather than a one-off effect.

5.2 RQ2. Can ATTR uncover valuable subgroups overlooked by rule-based targeting?

The second research question investigated whether ATTR could reveal hidden customer segments that rule-based approaches missed. WiDeX identified light users whose purchase frequency was only 1/100 that of premium subscribers yet whose viewing activity was 1.5× higher than traditional targets. ATTR converted this behavioral interest into a reliable targeting signal, producing conversion rates comparable to those of premium customers. Quantitatively, ATTR expanded the eligible target pool by 2.7× relative to rule-based campaigns and achieved a mean conversion uplift of 39.9 percentage points within this expanded segment. These values again represent normalized pp changes from each group’s baseline performance rather than absolute rates. Such improvement was not achieved by prediction alone. Explanations such as “consistent recent engagement with action films” provided counselors with persuasive entry points when approaching customers with little or no payment history. Monitoring logs showed a positive association between script usage and conversion rates, suggesting that operators adapted to evidence-based interactions over time. These findings demonstrate that ATTR enhanced both the discovery of neglected segments and the persuasiveness of operator-led campaigns.

Table 1. Comparison of rule-based targeting and ATTR targeting

Metric	Rule-based Targeting	ATTR Targeting
Pool Size	× 1	× 2.7
Avg. Purchases	× 1	× 0.01
Avg. Viewing Activity	× 1	× 1.78
CVR	× 1	× 1.1

5.3 RQ3. Is the observed uplift causally attributable to ATTR?

The third research question assessed whether the observed improvements could be causally attributed to ATTR rather than external factors. Because premium and frequent purchasers were reserved for existing rule-based logic, randomized A/B testing was infeasible. Instead, we employed a Difference-in-Differences (DID) design comparing conversion trajectories between the AI-targeted group and the historical control cohort of low-purchase users. Before the intervention, both groups displayed statistically parallel trends, validating the DID assumption. After the intervention, the AI-targeted group improved by 23.5 percentage points relative to its pre-period baseline, while the control group improved by 8.5 percentage points. The resulting net effect of 15 percentage points — computed on the same normalized pp scale — represents ATTR’s causal impact under enterprise constraints. Figure 2 visualizes these dynamics, highlighting parallel pre-campaign trends and post-campaign divergence. Although the study window was limited to four months, the DID framework mitigates major confounders and provides a robust approximation of causal impact in non-randomized settings.

6 Conclusion and Future Directions

This paper introduced ATTR, a practical experimentation framework for adaptable data-driven targeting under real-world constraints. Integrating recommender modeling, operator-ready explanations, and causal evaluation into a modular pipeline, ATTR links predictive modeling to measurable business outcomes. In practice, it raised conversion performance, broadened target coverage, and surfaced overlooked customer segments.

The explanation layer enhanced counselors’ persuasiveness, while causal analysis confirmed that improvements were driven by the framework rather than external factors. Though evaluated within IPTV marketing, the three-layer architecture is broadly applicable to other settings where randomization is difficult—such as telecom churn modeling, banking cross-sell, and media ad attribution—where credible causal evaluation remains crucial. For confidentiality reasons, implementation details are omitted; Section 3 provides a stepwise description of the pipeline for reproducibility. Future work will include releasing a synthetic dataset and module-level examples to support cross-domain validation.

REFERENCES

[1] Luciano Floridi and Josh Cowls. 2022. A unified framework of five principles for AI in society. Machine learning and the city: Applications in architecture and urban design. 535-545. DOI: https://doi.org/10.1002/9781119815075.ch45

[2] Bo Li, Peng Qi, Bo Liu, Shuai Di, Jingen Liu, Jiquan Pei, Jinfeng Yi, and Bowen Zhou. 2023. Trustworthy AI: From Principles to Practices. ACM Computing Surveys. 55(9). 1-46. DOI: https://doi.org/10.1145/3555803

[3] Nathalie A. Smuha. 2019. The EU approach to ethics guidelines for trustworthy artificial intelligence. Computer Law Review International 20(4). 97-106. DOI: https://doi.org/10.9785/cri-2019-200402

[4] Dominik Kowald, Sebastian Scher, Viktoria Pammer-Schindler, Peter Müllner, Kerstin Waxnegger, Lea Demelius, Angela Fessl, Maximilian Toller, Inti Gabriel Mendoza Estrada, Ilija Šimic´, Vedran Sabol, Andreas Trügler, Eduardo Veas, Roman Kern, Tomislav Nad, and Simone Kopeinik. 2024. Establishing and evaluating trustworthy AI: Overview and research challenges. Frontiers in Big Data. 7. DOI: https://doi.org/10.3389/fdata.2024.1467222

[5] Hugo Villamizar, Daniel Mendez, and Marcos Kalinowski. 2025. Towards a Framework for Operationalizing the Specification of Trustworthy AI Requirements. IEEE International Requirements Engineering Conference 2025 (RE '25). DOI: https://doi.org/10.48550/arXiv.2507.10228

[6] Yongfeng Zhang and Xu Chen. 2020. Explainable recommendation: A survey and new perspectives. Foundations and Trends in Information Retrieval 14, 1, 1–101. DOI: https://doi.org/10.1561/1500000066

[7] AKM Bahalul Haque, Najmul Islam, and Patrick Mikalef. 2023. Explainable Artificial Intelligence (XAI) from a user perspective: A synthesis of prior literature and problematizing avenues for future research. Technological Forecasting and Social Change. 18. DOI: https://doi.org/10.1016/j.techfore.2022.122120

[8] Koen W. De Bock, Kristof Coussement, Arno De Caigny, Roman Słowiński, Bart Baesens, Robert N. Boute, Tsan-Ming Choi, Dursun Delen, Mathias Kraus, Stefan Lessmann, Sebastián Maldonado, David Martens, María Óskarsdóttir, and Carla Vairetti. 2024. Explainable AI for Operational Research: A defining framework, methods, applications, and a research agenda. European Journal of Operational Research 317(2) 249–272. DOI: https://doi.org/10.1016/j.ejor.2023.09.026

[9] Léon Bottou, Jonas Peters, Joaquin Quiñonero-Candela, Denis Xavier Charles, David Maxwell Chickering, Elon Portugaly, Dipankar Ray, Patrice Simard, and Ed Snelson. 2013. Counterfactual reasoning and learning systems: the example of computational advertising. Journal of Machine Learning Research 14, 1, 3207–3260. DOI: https://dl.acm.org/doi/10.5555/2567709.2567766

[10] Stephen Bonner and Flavian Vasile. 2018. Causal embeddings for recommendation. In Proceedings of the 12th ACM Conference on Recommender Systems (RecSys ’18). ACM, New York, NY, 104–112. DOI: https://doi.org/10.1145/3240323.3240360

[11] Chen Gao, Yu Zheng, Wenjie Wang, Fuli Feng, Xiangnan He, and Yong Li. 2024. Causal inference in recommender systems: A survey and future directions. ACM Transactions on Information Systems 42.4. 1-32. DOI: https://doi.org/10.1145/3639048

[12] Huishi Luo, Fuzhen Zhuang, Ruobing Xie, Hengshu Zhu, Deqing Wang, Zhulin An, and Yongjun Xu. 2024. A survey on causal inference for recommendation. The Innovation, 5(2). DOI: https://doi.org/10.1016/j.xinn.2024.100590

[13] Jinmo Kang et al. 2024. A Novel Hybrid Architectures for Overcoming Sparse Data in Subscription Product Recommendation Systems. IEEE International Conference on Big Data (BigData '24). IEEE, 4104-4109. DOI: https://doi.org/10.1109/BigData62323.2024.10825732